is referred back to the facility used in (the illustrative application of) the CASTE equipment. Chapter 10 contains a discussion of broader, but actually realised, applications of the theory; mostly in education. Chapter 11 ties up the threads of the argument, some of them in a speculative fashion. We deal with the intimate relation between the (previously separate) macro-grain and micro-grain theories and, as a result, reach some conclusions about self-reference (the reflective theory) and the conditions for consciousness. In addition, we review some of the topics to be discussed in the next volume.

Materials that are not essential to the main line of the argument have been relegated to Appendices, in order to allow for the concise statement of fairly complex and closely knit ideas. The reader is urged, however, to regard the Appendices as part of the story; they are not just depositories for optional reading matter. For preference, they should be scanned when they are referenced, or shortly afterwards. If that is done, nothing will be lost by perusing them at leisure.





## 1. Observation and regulation

Any kind of measurement, regulation, or control calls for a standard experimental condition in which observations are made and parameters of the system are adjusted.

The classical paradigm was spawned in the physical sciences (though in that province it now has only the limited status of a paradigm for observing coarse grained phenomena). The standard condition is a fixed frame of reference, an observer of which is impartial and external to the experiment. In most cases he is omniscient; all possible outcomes (those that are countenanced as relevant to the enquiry) are known beforehand and the observer entertains a definite hypothesis, either deterministic or probabilistic in calibre. The system to be observed is demarcated unequivocably as an *it* under scrutiny and is often divided into parts such as transmitter and receiver or causative input and output effect.

Commonly, one of these parts, the *environment*, is designed by the experimenter/external observer to react in a particular way to the object/assembly under examination. Failing that, the experimenter acts himself as the environment but his actions, in this capacity, obey strict and predetermined rules. Apart from these interactions, usually represented as values of independent variables, all other fluctuations (deemed irrelevant) are minimised, i.e. the experimental conditions are held constant, at any rate to the extent that they might influence the values of observed or dependent variables. If the experiment is repeated, every effort is made to replicate the same conditions. Finally, for several reasons (one of them the reductionist ethos underlying this method, another the ease of replicating particular entities) there is a tendency to carry out experiments that involve discrete perturbations (or the variation of one independent variable at once) and the observation of simple reactions (or few dependent variables).

1.1. *Behaviouristic formulation.* The classical paradigm is carried over, more or less piece-meal, into behaviouristic psychology though the nomenclature is changed. For example, to secure constant conditions, a rat is starved prior to entering a maze in which food will be delivered as a reinforcement; the maze itself is stripped of olfactory or tactile cues whereby "irrelevant" features of junctions might be recognised. The hypothesis, in this case, is that rats escape from mazes towards food, that the alternative (relevant) behaviours are movements, and that the animal adapts after repeated runs. Any specific hypothesis (whether it pertains to a rate of adaptation, a likely path, or a complex intermediary notion such as approach/avoidance conditioning) is rooted in these general assumptions and others of the same ilk. The object or assembly under scrutiny, given the generic title of *organism* in this area of science, is the rat. The environment is the maze. It might equally be a piece of preprogrammed equipment; for example, the control system for a Skinner Box. More obviously in the latter case (but in fact, in the maze example also) observed responses are simple movements, regarded as reactions to stimuli that are either discrete and simple in themselves, or decomposible into discrete and simple elements.

Even at this stage it is clear that the meaning of terms like "parameter" and "stimulus" and "response" becomes very strained if the experiment involves a human being. The behaviouristic version of the classical paradigm, as described a moment ago, has a certain reality and a certain face validity if the organism is a rat or a pigeon or a particular reflex subsystem of a man. But consider what happens if the organism is a human being. Starvation, for example, is out of court (for psychological work, that is; physiologists may legitimately starve their subjects and even implant electrodes in them). With rare exceptions (sensory deprivation in the adult, deprivation of peer contact, affection, and so forth, in child psychology) the experimenter issues instructions in the belief that they will induce an appropriate mental "set". A sizeable body of literature is devoted to discussion of what instructions can be delivered to replicate the conditions on different occasions and for different people. The trouble is that instructions establish constancy (if they do so at all) because they are interpreted, understood, and accepted; not because they are imposed, like a quasi physiological constraint such as "raised hunger drive". Because of that, instruction giving and instruction accepting is outside the classical framework; it belongs to a normative situation, not to a functional situation bounded by conditions the experimenter is able to set up independently of the



organism. Actually, very similar comments apply to the exceptional case noted a moment ago, for example, sensory deprivation and deprivation of peer contact act as symbolic operations that are understood in terms of security, fear, status, etc. Hardly ever do they act simply as surrogates for the kind of change that may be induced (for instance) by simulating the head of the caudate nucleus (one "negative" centre), or administering the Toad Poison, Bufotonin.

1.2. *Criticism of the classical (behavioural) paradigm.* Although the classical paradigm does provide an acceptable framework for gross physiological observations (at the level of fine grained observation it is open to the same criticisms as the corresponding construct in physics) there are reasons to believe it is literally irrelevant to psychology. Several basic objections become evident in the course of the discussion; the strongest being that this experimental framework is unsuited for the observation of conscious phenomena which (according to the present point of view) form the principal subject matter of psychology. Further, the paradigm forces a false dichotomy upon the experimenter. It leads to the counterpositioning of "subjective" and "objective" observations as though the former were fallible (though inherently interesting) and the latter concrete and reliable though often (as Von Foerster, 1971, demonstrates) uninformative. In fact, the experimental structure prescribed by the classical approach simply excludes the observation of subjective events, by edict, and permits only a particular sort of objective observation. The constraint springs from the requirement that any organism is regarded as an it, and has no bearing on whether subjective events, if they were observed, could be reliably observed or not. Because of this exclusion, it turns out that any "subjective observation" is consigned to an other-than-standard observation scheme; usually introspection or retrospection of which the hazards, though probably overrated, are acknowledged. We seek, and will submit, an experimental scheme that will accommodate subjective events (though by no means usually introspections) as the minimal frame of a psychological experiment, or on a broader scale, the least structure in which it is possible to control or regulate a psychological process of any kind whatsoever. Attention is focussed upon the processes *learning* and *teaching* and *cognising* but other specific mental activities receive detailed consideration; notably *examining* and *model making* and *innovating*.

1.3. *Normative paradigm.* Insofar as this criticism holds good, it would be better to admit the normative and symbolic character of the situation at the outset. In this case "giving instructions" simply glosses the act of "establishing an experimental contract in which the subject agrees to act in a given fashion and to speak in a certain language". As it stand, the subject, not the experimenter, is held almost completely responsible for maintaining the conditions under which the experiment is conducted. The experimenter usually gives minimal assistance in this matter; for example, by providing a soundproof room or an evenly illuminated and uniformally textured background (superficially, these acts are analogous to removing irrelevant junction cues in the rat maze; in practice they serve primarily as conditions which allow the subject to obey the instructions if he is willing to do so and thus play a distinct role).

There is nothing especially outlandish about a normative framework; such schemes are used frequently in anthropology, social and political science and various other fields of enquiry; their relation to other schemes is discussed more comprehensively in Pask (1966, 1971a). Perhaps the most apposite example, in the context of learning and teaching (and game playing which is an important concomitant of both) is given in *Homo Ludens* (Huizinga 1949) and the specially sterilised encapsulation of these ideas, Game Theory, is shown in Chapter 7 to have immediate relevance.

The frank admission that human experiments or human control situations are set in a normative framework, rather than the functional or quasi physical framework of the classical paradigm avoids the extravagant and odd usage of words like "stimulus" and "response" which have a perfectly good meaning in the context either of physiology (stimulus $\triangleq$ effective stimulus) or the semantically neutral terminology of system theory (stimulus $\triangleq$ causative input at a given instant). Clearly, anything the experimenter does to the organism may be regarded as a stimulus (for example, reading a poem or showing a film, giving a command or asking a question) and any correlated action may be viewed as a response. The classical experimenter is bound to view it thus, for exactly this meaning is induced by the experimental paradigm. But this is not a matter of fact. Other realities are just as legitimate; for example, the poem/film may be regarded as a poem/film that the organism understands; the command as a directive to be obeyed,



the question, as a comprehensible enquiry. None of these are only stimuli, none of them need be viewed as stimuli at all. Similar comments apply to the organisms actions or replies or explanations (the word "response" is less loaded with extraneous meaning; a reply or an explanation is a response; but it is not a reaction to a question).

1.4. *Diversity of rules.* In other words, the experimenter/observer can stick to many kinds of rule. If he wants to entertain the hypothesis that an organism (*it*) will respond to a stimulus, that is quite legitimate and quite on a par with "the organism reacts with a reflex knee jerk to a tap on the patellar tendon". On the other hand if he adopts a normative and language-based framework then he can give the organism commands and ask it questions in a language that is part and parcel of the normative scheme. In this case, however, there is a sense (to be developed at some length) in which the organism, qua sentient being, is no longer "it". Either set of rules is acceptable (amongst other variants and possibilities (Pask 1971b, 1972)); only misusage is to be deprecated. It is logically incontinent to switch meanings between the classical paradigm and any variety of normative paradigm. Talking about "attention-direction stimuli" and "questioning stimuli" and "hypothesis-indicating stimuli", as though the corresponding commands, questions and requests for explanation are only (very complex) stimuli, is sheer twaddle.

These criticisms admittedly smack of pedantry in the context of experiments (such as the majority of perceptual experiments) in which situations are broken down into discrete elements. It is generally true that these experiments can be treated as nearly physiological. That is also true of experiments upon adaptation, habituation, and conditioning. But the criticisms are not at all pernickety in the context of learning; especially when a human being is learning to perform a complex perceptual motor skill or learning to recall a sensibly large body of subject matter. This is the area for which the present theory of cognition and learning is primarily tailored and to which it is most profitably applied. As the discussion proceeds it will become evident that a fully fledged normative interpretation is mandatory and that a theory of learning, constructed within such a framework, involves, as a dual, a theory of teaching (conversely, a theory of teaching carries with it a theory of learning).

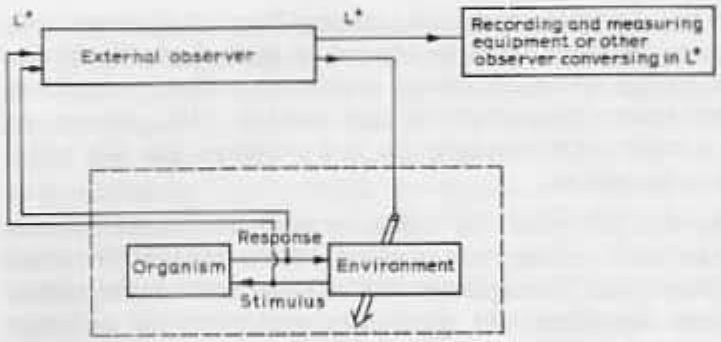


Fig. 1.1. Status of external observer in classical situation.

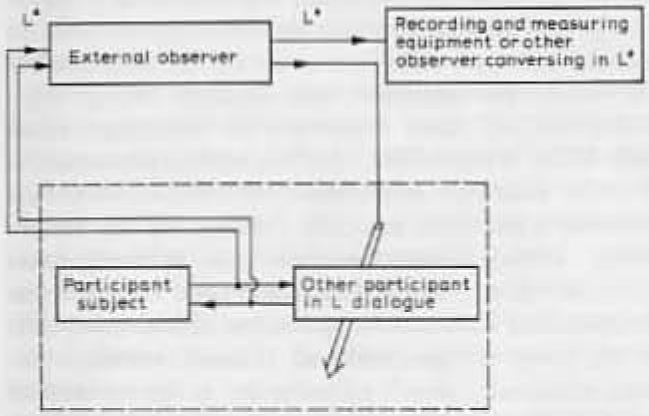


Fig. 1.2. Status of external observer in normative situation.

1.5. *Summary.* To survey these notions, the status of an external observer in the classical paradigm is shown in Fig. 1 and his status in the (relevant) normative scheme is shown in Fig. 2. The two picture forms are identical (hence a fairly widespread but mistaken equation of the two paradigms), but the labels are utterly different.

1.5.1. In Fig. 1 the experimenter is able to distinguish an organism and its environment on grounds that are independent of the nature of a subsequently observed interaction; for example, by placing a spatial boundary around each one. Either in reality (the experimental scheduler for a Skinner Box) or in effect (the experimenter follows a rigid algorithm and delivers a schedule of



stimuli, reinforcers, and the like) the environment reacts according to a specific formula; so far as the external observer is concerned, it acts causally in respect of the organism. That is, it delivers stimuli (perhaps contingent upon previous events) and for each stimulus the observer entertains an hypothesis (perhaps probabilistic) which has been built into the environmental formula, regarding the possible responses of the organism. The events which occur are described, in terms chosen by the external observer, as a behaviour consisting in stimuli (maybe reinforcers) and responses and the environment is designed to react in terms of this externally chosen description. The description itself, together with the operating principles of the classical situation is described in terms of an observational metalanguage for talking about experiments, designated L*.

It is quite possible for the external observer to change the characteristics of the environment from time to time and he may do so, for example, when alternating experimental treatments. This act is represented by a parametric arrow passing through the environment box. As a further refinement, such transformations of the formula in the environment may be monitored, by a sequential sampling scheme; for example, on the basis of observations of the stimulus response interaction between the organism and its environment. If so, the possible formulae are known beforehand and, in principle, the experimental strategy might have been built into the environment itself. The experimental objective is generally taken to be the accumulation of enough stimulus/response or input/output data to characterise the organism either (in computational terms) as a finite automaton or (in continuous mathematics) as a transfer function. A typical control objective is to design an environment so that a performance index (some function of the behavioural history) is maximised; for example, to choose a reinforcement schedule that maximises adaptation rate.

1.5.1.1. Two unobtrusive aspects of this picture are theoretically crucial; the observer's choice of a description and the limited scope of his parametric operations. Between them, they determine the environment as a finite automaton (and the observer's hypothesis as a hypothesis under the adumbrating dogma that the organism is to be characterised as a finite automaton). The fixed description determines the input states and output states of these automata and the parametric restriction implies that the internal state set of the environment cannot be changed as a consequence of any act (like attending to something novel) that the organism might perform outside the limits of the initial description. Given this, it is possible to apply a very elegant argument due to Von Foerster, 1971, (as part of this theory of finite functional machines) in order to show that any behavioural sequence produced by the interaction in Fig. 1 is, in a sense, predetermined. For example, the proposition that the coupled system "organism/environment" will undergo adaptation is tautologous; the system must do so unless the entire paradigm is misapplied (there would be some condition that is not held constant valued as required by the classical paradigm). Further, all modes of adaptation are predictable and none of them, in the sense of the theory to be discussed, constitutes learning.

1.5.1.2. Due to the automaton theoretic identification, an observer is able to interpret the interaction between the organism and the environment as an exchange of symbols in a formal language. This point is important because the word "language", though used in accordance with automaton theory and mathematical psycholinguistics, neither stands for language in the ordinary sense nor for language as we use the word, generally, in this book. The formal "language" in question is a syntactic entity only. The reader familiar with automaton theory will notice that the organism and the environment may be regarded, equivalently, as abstract "devices" for accepting/generating strings of symbols; that the alphabet of this formal "language" consists in description tags for events called (in L*) "stimuli" and "responses"; that these symbols label state transitions (input accepting and output generating); that the "internal" states of the automata representing the organism and its environment are "non terminal symbols" in the formal language and that the structures of these automata may be regarded as formal "grammars" or ordered collections of rewriting rules.

1.5.1.3. Looking on from outside the system, an external observer is able to form contingency tables depicting symbol or symbol-string occurrences; to tabulate combinations of stimuli and responses. By forming ratios of the frequencies with which combinations of events occur, it is possible to estimate, in one way, event probabilities, conjoint event probabilities, and so on. From this data, the external observer may, if he wishes, construct indices of uncertainty from his point of view and of information



(uncertainty reduction) also from his point of view. Using the calculus of information measures (transmissions, interactions and the like) the external observer can construct an elegant macrotheory of what goes on.

The operation is chiefly of consequence because this usage of macrotheory does not exactly tally with the usage exemplified later in this book. Confusion is possible for the naming is often identical (we also talk of information and uncertainty) and the calculus is identical. But it is important to recall, when this point is reached, that the information/uncertainty indices of Chapter 2 Section 4 estimate doubt or belief entertained by the subject and not, in general, by the external observer.

1.5.2. In Fig. 2 (as in Fig. 1) the observer uses a metalanguage L* in order to describe a system, to prescribe actions, to pose and test hypotheses that he may discuss with other observers. However, the status of the two lower boxes is entirely different. They represent participants and they are distinguished as normative entities, by dint of an experimental contract which is negotiated in preliminary dialogue with at least one of the participants (conventionally, the left hand box) who is a human being. At one extreme, the right hand box may also represent a human being; at the other extreme, it may be a finite state machine called the environment (as in Fig. 1); several intermediary cases are also of great interest.

1.5.2.1. Suppose the right hand box is an environment; structurally represented as a finite state machine. The contract to act in a particular manner is established with the left hand participant only. The preliminary dialogue used to negotiate the contract (in the limit, to give instructions) is either L* or some other tongue with most of the capabilities of a natural language. The tenure of this contract, throughout the experiment, depends entirely upon the participant's cooperation. Although the interchange between this participant and the environment is essentially a stimulus response interchange, the onus for keeping it so (i.e. for keeping an agreement to act like an automaton) rests on the participant's shoulders and it is convenient to think of a special cognitive compartment in the participant set aside for this purpose. The mental activities going on in this special compartment are private; they are rendered private because, in this extreme case, all that an observer is able to examine (namely, the interaction between the boxes, however they are specified) has been restricted to a form which is unable to represent the maintenance of a mental "set". By edict (and by the contract that embraces and reifies the edict) the participant acts causatively on the environment and vice versa.

Of course, the observer could examine the maintenance process after the experiment by eliciting retrospective commentaries (in L*) from the participant, or he might sample the process throughout the experiment by a protocol (which is tantamount to embellishing the system with a further quasi natural language). Both of these expedients are open to valid criticism; retrospection on the grounds that recall often distorts a message and introspection, by protocol, on the grounds that it interferes with the experiment (if an observer accepts introspections from a participant he is no longer an external observer and no longer regards the participant as an impersonally pronominalised object of observation).

1.5.2.2. At the other extreme, both the lower boxes in Fig. 2 are interpreted as participants and they are in conversation. Their interaction is not, in this case, simply a stimulus/response interchange but is dialogue in a language, L, which is called an object language, to distinguish it from the observer's metalanguage L*. As part of the contract the participants agree to use L throughout the experiment. L need not be a spoken or written natural language (a film was deliberately cited earlier as well as a poem; in general, L transactions can take place in any modality and may even be mediated by mechanical operations like button pressing). But L must have some of the qualities of a natural language and though formalisable (incompletely perhaps) it is not a "formal language" of the type employed to represent the "stimulus response interchange" of Fig. 1. At least L must have a pragmatic and a semantic interpretation, over and above its syntax; it must be a command and question language and it must admit ostension and predication.

At this extreme, the responsibility for keeping the original contract is distributed between the participants and they may employ L expressions in order to discuss the matter, or, in general, to externalise normally private mental operations as stretches of L dialogue. An external observer can, as before, record these transactions.



The disadvantages of the arrangement are as follows: (a) It is not obvious that a conversation will take place. (b) Even if it does (so that the observer has dialogue to record) he may be unable to interpret the dialogue unambiguously. (c) A participant is an organisational (or normative) entity; as such "participant" is well defined. But it does not follow that the organisational participant (which, in this arrangement is the meaning of a box) has a fixed spatial location; for example, in one person's head.

Obviously (a), (b) and (c) are difficulties that beset the observation of any conversation, whether or not it takes place in an experimental or control situation. To some extent the intermediate constructions, mentioned earlier, overcome these difficulties.

1.5.2.3. An intermediate interpretation is given to Fig. 2 if the external observer specifies or partially specifies one of the participants as an *agent*. He could do so, for example, by instructing a human *interviewer* or a human *teacher*; both of them roles that to some extent pervert "free" conversation (making it an interview or a tutorial). Equally, instead of a human agent, the observer might introduce a mechanical processor (not exactly a finite state machine) and make this engine execute his instructions as a participant. In the sequel, special attention is given to two (humanly or mechanically executed) designs of this type.

(1) An experimental contract is established whereby the participating subject aims for a goal which he cannot actually achieve on his own. The observer's participant gives the subject the cooperative assistance needed in order to satisfy the experimental contract if and only if he engages in dialogue, and by means of it, externalises the (normally private) cognitive events involved in keeping his contract or, equivalently, satisfying the agreed goal. This method is called a *cooperative externalisation technique* (or CET) and the series of instructions characterising the observer's participant (whether executed by a human being or a machine) is a CET heuristic.

(2) The other design incorporates a specialisation or restriction of the CET heuristic and has several variants more profitably detailed in context. In all of these variants the observer's participant acts as a conversationlist who is a teacher (in a weaker or stronger sense but never to the extent of rendering the underlying CET heuristic inoperative). The observer's participant is characterised as a *tutorial heuristic*.

Let us start anew on a slightly different tack (very soon the lines of argument will converge upon a unified picture of a relativistic psychology).

The fresh tack lays emphasis upon the dynamic properties of a human being; perhaps of any candidate "organism" an external observer opts to consider. Most theories, behaviouristic or otherwise, credit the organism with a more than passive role. In addition to reacting, with respect to an input or stimulus for example, the organism is supposed to produce and emit autonomous responses or operant responses. For instance, the whole discipline of behaviour shaping depends upon the production of operants that can be modified by an appropriate choice of contingencies. It matters little whether the autonomous form of output is ascribed to an "internal stimulus" that produces it (unlike an "attention directing stimulus", so violently criticised earlier, an "internal stimulus" is a perfectly legitimate construct; it has the same status as any other stimulus and no essential connotation is glossed over).

2.1. The existence of a stream of autonomous responses occurring at a certain average rate, places constraints upon the experimental structure that are generally respected in practice but omitted from the theory. If there is a certain field (stimulus response set in the classical case; a field of attention in a cognitive image of things) then these responses must do something; at least they "stimulate" or "provide an input to" the organism. Hence, if an external observer is anxious to encompass the entire field of an organism or subject the field he chooses must have a certain minimal size. The autonomous responses will be produced, willy nilly, and he will be unable to regulate their production unless the environment he designs does something about each one; for the rest of them will go unnoticed. The problem is far more obtrusive in case the external observer wishes to regulate a process rather than observe it and in this connection any environment that purports to be a regulator must satisfy Ashby's (1964, 1st edn. 1956) "law of requisite variety"; it must be equipped with at least the "variety" required to counter the variety of autonomous outputs.

In human terms, a subject must be given enough to do. Moreover, if we want the subject to focus upon an environment



we deem relevant then the subject must be given enough relevant things to do. Failing that, he will succumb to tedium or boredom and do irrelevant things; and these, in general, will be concealed from view.

There is, of course, a converse limit which receives more discussion, especially in connection with the psychology of skill. Suppose that an environment (for men, a task to be performed during an experiment) provides a sequence of inputs in excess of a certain limit, then the subject is overloaded and fails to perform the task at the requisite pace; either he may make errors (when, at a slower pace, no error would have been committed) or he may omit certain facets of the task or he may simply give up and do something different. Since man at least adapts (a truism, for any large dynamic system does so) there is no genuine "channel capacity" in Shannon's (1949) sense. For the load imposed by an input may depend upon all of the inputs received up to the moment in question. Or put commonsensically, as it may be if the subject is assumed to learn input/output relations, a once difficult task will become less difficult as the task relation is learned. However, taking account of this tendency and conjoining the underload and the overload limits it makes sense to conceive of an operating region in which it is possible for the subject to exist in equilibrium with an environment he has agreed to focus attention upon. Within this operating region there is enough to do but not too much to do.

Generally, it is hard to give clear meaning to "too much" and "too little". The matter becomes more tractable if the "things to be done" are conceived as "problems to be solved" (rather than inputs and outputs of an automaton) for, in this case it is possible to posit an index of difficulty to the subject or conversely an index of problem simplification, as well as an index of performance.

2.2. Suppose, until the matter is questioned, that a difficulty index exists and is given the generic label $\eta$. Suppose that the increments in value of $\eta$ are (in a sense to be examined) small enough. Suppose, also, a source of problems which can be displayed to the subject, and (as a condition which ensures that the subject can be overloaded by a relevant task) suppose the most difficult problems, indexed by the maximum value of $\eta$, are too difficult for the subject to solve. The source may either be associated with an index $\eta$, so that problems of an arbitrary difficulty can be selected or, more conveniently, suppose the source emits problems of maximum difficulty ($\eta_{max}$) and that there are simplifying operations, indexed by $\mu = \eta_{max} - \eta$, which are applied to simplify or partially solve the problems selected by the source before they are displayed to the subject. Let $\rho$ be a performance index such as correct response frequency or latency weighted correct response frequency or a decreasing function of the error modulus integral over a certain past time or of the RMS error. The appropriate performance index, $\rho$, and the physical variable corresponding to $\eta$ depends upon the kind of problem solving task.

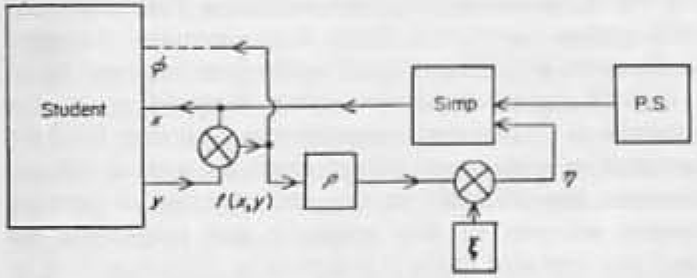


Fig. 1.3. Stimulus display = x; response = y; "correct" comparator output = f(x,y); P.S. = Problem Source; Simp = Problem simplification or difficulty variation; $\phi$ = Knowledge of results, if given; $\otimes$ = comparator symbol.

2.3. A steady state control system (Fig. 3) is a device that interacts with a subject who is performing a task in order to maintain his proficiency at a near constant level by posing problems at an appropriately selected level of difficulty, $\eta$. If $\rho$ increases $\eta$ is made to increase and vice versa so that $\rho$ approximates a stipulated constant value $\xi$. The controller has the (ideal) effect of clamping the subject in a working region where the problems posed are neither so difficult as to prove incomprehensible nor so easy that there is insufficient relevant variation to occupy his attention. By prior hypothesis, either circumstance would lead to a state of affairs in which the subject would be bound to focus his attention upon problems deemed irrelevant by the external observer (a subject can neither assimilate problems that he fails to comprehend as problems, nor can he stop solving problems; it is only a matter of whether he solves relevant



problems). Viewed thus, the steady state controller is an instrument that captures the subject's attention or, more accurately, that permits him to honour an agreement to solve problems of a particular kind.

Controllers either of the type shown in Fig. 3 or the refined versions that are described below, have been studied empirically in connection with many different tasks; for example, pursuit and compensatory tracking, visual discrimination, trajectory detection, logical puzzle solving and operating typewriters, card punches and other keyboard instruments, (Gaines 1967, 1972; Hudson 1962; Kelley 1968, 1970; Pask 1966).

Some typical identifications of Fig. 3 for a tracking task, are as follows. The performance index, $\rho$, is the converse of an RMS error, measured from target (pursuit tracking) or a null position (compensatory tracking). The problem is to catch up with a haphazardly moving target (pursuit) or to maintain a vehicle stable under the influence of a perturbing input (compensatory tracking). The problem difficulty, $\eta$, for both pursuit and compensatory tracking may be varied either (a) by increasing the mean amplitude of a random forcing input to target or vehicle, (b) by operating upon the vehicle simulator, which relates the variable (usually acceleration) manipulated by the subject to the displayed vehicle position, (c) by delivering auxiliary "cueing" information for example, about the vehicle's velocity, (d) by introducing or withholding a requirement to perform an auxiliary task.

2.4. Under conditions that satisfy the initial assumption the conjoint steady state system, made up from the subject and the controller, is empirically stable; that is $\rho$ does approximate $\xi$ by dint of the prescribed variation. More generally, for the multi-dimensional systems considered below, the joint system is Lyapunov stable. That is, the time derivative(s) of $\rho$ (or all terms $\rho_i$; i an index over the dimensions) may be set to zero to express the system equations in terms of derivatives of $\eta_i$ (or $\mu_i$). On representing these quantities by a cyclogram in coordinates $\mu_i$ it is possible to plot Lyapunov functions as shells at varying distance from the origin. "The system is stable" means "that its state trajectory approaches the origin by cutting through these shells inwards and never traversing them in an outward going direction".

The method of securing stability by this means is called the *steady state technique* or, for brevity, the SST.

The SST has various practical applications of which the following are representative though not exhaustive.

(a) As a psychological measuring technique (for example, in aptitude testing). The basic index is the amount of compensation needed in order to secure a standard level of performance or to do so over a fixed interval (Gaines 1969, Kelly 1971, Pask 1971c).

(b) As a device for maintaining vigilance and attention (or determining the effects of fatigue). In this case "false" alarm signals are injected into the main signal channel and the subject's response is ascertained. If he misses "false" alarm signals the interpolation rate is increased, sometimes differentially for distinct classes of signal (Pask 1960, 1971).

(c) Given an intermittent "interfering" task and an SST regulated main task, as an indicator (in terms of the compensatory $\xi$ level with and without the interfering task) of either the subject's division of attention or the interfering task's distraction value (Gaines 1972, Pask 1969, 1970).

(d) As a means of bringing about standard psychological conditions other than constant loading; for example, a constant level of ambiguity or self consistency in judgment (Kelly 1968, Pask 1964).

(e) As a dynamic analogue to the "staircase" method in psychophysics or to "titrate" a physiological variate such as integrated EMG against any appropriate task variations (Verhave 1963, Kelly and Prosin 1968, Pask 1971c).

2.5. Scrutiny of the time trend in any system trajectory (a typical example is shown in Fig. 4) reveals a general increase in the value of $\eta$ needed to obtain and maintain the condition $\rho = \xi$ as time progresses. The trend is due to the fact that a subject learns the task relation(s) and thus needs to deal with more complex


Fig. 1.4. Response over trial intervals (t) of steady state system: shown in terms of "difficulty" (——, $\eta(t)$) and proficiency (---, $\rho(t)$).



problems if he is to remain within his own operating region. This leads naturally to another and fully compatible interpretation of the steady state technique; namely, that the SST compensates for the influence of learning; as the student learns, so he is given more to learn about. One interesting consequence of this interpretation, from the external observer's point of view, is the reduction of adaptation to control. Thus, if viewed alone, the subject is an adaptive system (because, as a matter of fact, he learns a task relation). But the coupled man-machine system is a plain control system, in which the adaptive process is embedded. Another consequence is that stability cannot be maintained indefinitely; any device has a limited range of compensation at the extremity of which it is presenting maximally difficult problems; beyond this range, it is unable to obey its design rule; which is, incidentally, one straightforward application of Ashby's (1964) law of requisite variety.

It is also true that adaptation rate is maximised for a suitably chosen value of $\xi$ and because of this the SST has an incidental teaching or training function. For example, if $\rho$ scaled to the interval 1, 0, is a correct response frequency or a latency weighted correct response frequency or a sequential response function, there are many skills for which $\xi$ values in the region of 0.75 are "optimal" and fairly sharply defined with limits between 0.6 and 0.8 (Fig. 5). However, individual differences exist and even for superficially "homogenous" tasks like compensatory tracking with velocity or acceleration control, the optimal value of $\xi$ is prone to change as learning proceeds.

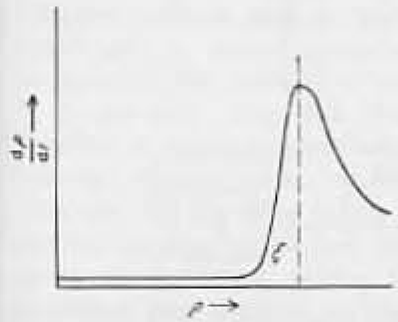

Fig. 1.5. Sketch showing form of "learning rate" as function of operating proficiency level in the imaginary case of a steady state system not operating; its operation holds $\rho = \xi$.

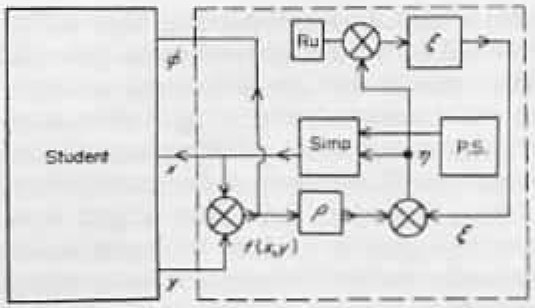


Fig. 1.6. x, y, f(x,y), $\phi$, Simp, P.S. as before; Ru is adaptation rule.

An adaptive teaching system (Fig. 6) takes account of these fluctuations by adjusting the value of $\xi$, usually by a simple "hill climbing" method, to satisfy an overall criterion of performance (for example, to maximise the product $\eta \cdot \rho$ at each sample point). More elaborate adaptive systems consist in several subcontrollers and an overall regulator (Fig. 7). These systems operate multi-

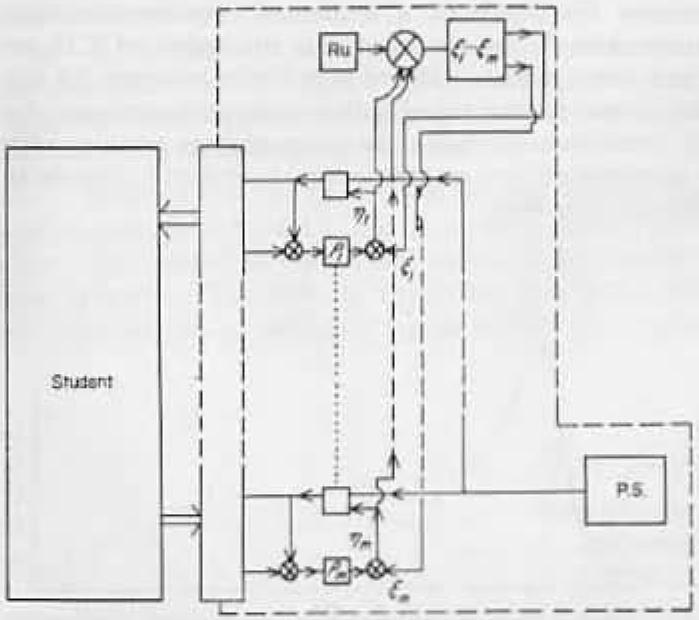


Fig. 1.7. Multidimensional adaptive system. Note single problem source (P.S.). System has m subcontrollers and one overall controller.



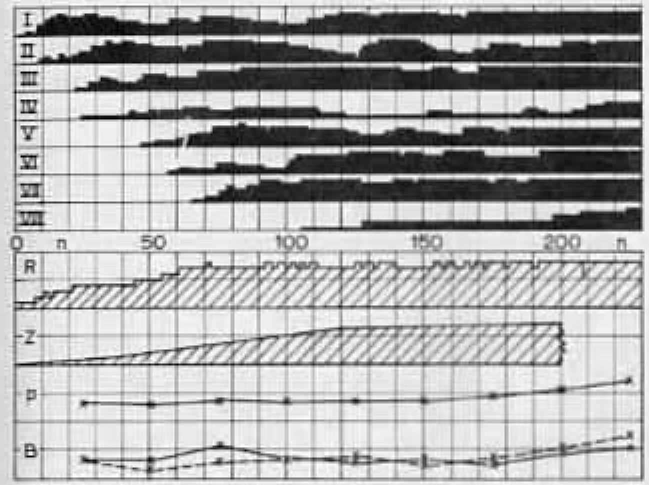


Fig. 1.8. Characteristics of a coupled system used for instructing a structured discrimination skill having eight different problem types. n = Number of trial; I. . .VIII = $\eta_I(n) \ldots \eta_{VIII}(n)$. R = Number of elements in subset from which subskills are selected. Z = Sum of the averages of the eight different subskills. $\rho$ = Proficiency measure defined as $\rho(n) = (1/R)\ \Sigma_i\ \rho_i(n)$. B = Proficiency measure with reference to rapid and less rapid predictions (See Pask, 1965b).

dimensionally; the several subcontrollers indexed i, maintain or approximate $\rho_i = \xi_i$; by adjusting difficulty variables $\eta_i$, the overall regulator adjusts the $\xi_i$ to maximise some function of the $\eta_i$ (for example, the sampled product of their values) and the stable trajectories have the form shown in Fig. 8.

2.6. Many of the conclusions about the behaviour of SST and adaptive systems used as teaching and training devices have been confirmed in recent years and appear to be quite commonly accepted in the context of computer aided learning. For example, it is generally conceded that feedback (to be effective) must confirm or deny an hypothesis that the student has formulated on his own account; that feedback does not just act to reinforce a response. By the same token task difficulty must be calculated in respect of a particular task description and the level of difficulty must be regulated within this framework, as a function of individual performance, so that problems are neither unintelligible (at one extreme) nor too easy to solve (at the other). The lag in reaching these conclusions is not altogether surprising. The experimental regulators employed over a decade ago were far more adaptive and complex than the majority of computer aided learning systems; at least this level of complexity is needed to obtain the results in question or even to reify the hypotheses. Of course, these regulators were built for individual use; in contrast, designers of computer aided learning systems work under additional constraints such as providing an organisation able to service a number of student consoles or terminals. That is the main reason why the adaptive sophistication of these devices has only just caught up with the sophistication of much earlier experimental machines.

The recent work is voluminous and widely dispersed. Some of it is reviewed in Pask (1970, 1971); some noteworthy progress in Great Britain, since that date, is due to Hartley and Sleeman at Leeds. Appendix A contains a summary of our own work on the tutorial application of SST and adaptive systems. The summary is focussed upon the most comprehensive studies (Lewis and Pask 1965, Pask and Lewis 1966, 1968) which refer to the accquisition of complex coding skills (these are perceptual motor skills with a substantial intellectual content as well). Appendix B is a similarly brief outline of studies, in the same context, of tutorial applications of the cooperative externalisation technique or CET. Experimental details of these and other systems are given in Pask (1973).

2.7. As an overall conclusion, the SST realises a standard condition for measurement which is quite distinct from standard conditions obtained by holding all but a relevant variable constant. The standard in question is a standard operation region or working region in which a subject is able to act as a problem solver/learner with respect to whatever the external observer deems relevant. This condition is secured by introducing a further entity, the controller, which has the status of a participant and performs either measuring or regulating operations relative to whatever the individual subject's characteristics may be. The kinds of compensatory variation so far discussed are quite elementary but do, nevertheless, exhibit the gist of a relativistic psychology; that is, of a psychology in which the external observer views the subject as a participant relative to some other participant in the context of a task or occupation. Insofar as human beings learn (and it will be argued that, whatever else may be true, human beings are designed to or impelled to or built to learn), the notion of a dynamic standard condition (for measurement or further control of a mental process) goes hand in hand with the notion of psychological relativism. For example, in the adaptive SST system, the adaptation rule is executed (to maximise learning rate) in the context of a dynamic-equilibrial or standard condition maintained by one or more SST subcontrollers.



2.8. Recall that the (theoretical) stability of any SST system depends upon the cluster of assumptions listed in Section 1.5; the empirical stability is, of course, just a matter of fact. Even for elementary modes of SST operation, these assumptions can be divided into two groups (both of which prescribe conditions that must hold). The first group characterises the subject as a participant in an albeit primitive conversation; hence, as an entity which is appropriately modelled within one or other form of the normative framework shown in Fig. 2 and unamenable to representation within the purely functional framework of Fig. 1. In particular, these conditions apply to the stable operation of any cooperative externalisation technique, or CET (Section 1.5.2.3) and to any tutorial mode derived from it.

The second group of assumptions are to do with special conditions on the task and on a subject's cognition that must be satisfied if a closed and simple minded participant like the SST controller is to maintain "discourse" (and, in the process, stability). These conditions are very restrictive and it is entirely possible to relax them to obtain non-trivial conversations which, being relativistic systems like an SST system, extend the scope and also the meaning of a relativistic psychology. In particular, a realisation of the CET heuristic or the tutorial heuristic of Section 1.5.2.3 is a relativistic system (of participants).

For the most part, the subsequent discussion is confined to conversations that highlight learning and have educational significance. In this context the main tenet of relativism was previewed at the end of Section 1.4; there is no theory of learning apart from a theory of teaching; no theory of teaching apart from a theory of learning. There is a theory of learning and teaching together and that is all. Results from many studies support this point of view as, also, does the evidence of commonsense. So, to be dogmatic (but with confidence) learning implies teaching and teaching implies learning. Sometimes, the teacher and learner responsible for the joint process are obvious (a student at a desk and another person wearing an academic cap and gown). Sometimes, the teacher and learner are not so obviously distinct and turn out to be unexpected but, once-indicated, intuitively plausible entities.